arXiv:1611.05911v1 [cs.DS] 17 Nov 2016

Computing Absolutely Normal Numbers in Nearly Linear Time

(Preliminary version)

Jack Lutz* Elvira Mayordomo†

November 21, 2016

Abstract

A real number x is *absolutely normal* if, for every base $b \geq 2$, every two equally long strings of digits appear with equal asymptotic frequency in the base-b expansion of x. This paper presents an explicit algorithm that generates the binary expansion of an absolutely normal number x, with the nth bit of x appearing after $n\text{polylog}(n)$ computation steps. This speed is achieved by simultaneously computing and diagonalizing against a martingale that incorporates Lempel-Ziv parsing algorithms in all bases.

Keywords: algorithms, Lempel-Ziv parsing, martingales, normal numbers

1 Introduction

In 1909 Borel [4] defined a real number α to be *normal* in base b ($b \geq 2$) if, for every $m \geq 1$ and every length-m sequence w of base-b digits, the asymptotic, empirical frequency of w in the base-b expansion of α is b^{-m}. Borel defined α to be *absolutely normal* if it is normal in every base $b \geq 2$. (This clearly anticipated the fact, proven a half-century later, that a real number may be normal in one base but not in another [8, 22].) Borel's proof that *almost every* real number (i.e., every real number outside a set of Lebesgue measure 0) is absolutely normal was an important milestone in the prehistory of Kolmogorov's development of the rigorous, measure-theoretic foundations of probability theory [16]. For example, it is section 1 of Billingsley's influential textbook [3]. The recent book [7] provides a good exposition of the many aspects of current research on normal numbers.

Borel's proof shows that absolutely normal numbers are commonplace, i.e., that a "randomly chosen" real number is absolutely normal with probability 1.

*Department of Computer Science, Iowa State University, Ames, IA 50011 USA. lutz@cs.iastate.edu.

†Departamento de Informática e Ingeniería de Sistemas, Instituto de Investigación en Ingeniería de Aragón, Universidad de Zaragoza, 50018 Zaragoza, SPAIN. elvira@unizar.es.

Rational numbers cannot be normal in even a single base b, since their base-b expansions are eventually periodic, but computer analyses of the expansions of π, e, $\sqrt{2}$, $\ln 2$, and other irrational numbers that arise in common mathematical practice suggest that these numbers are absolutely normal [5]. Nevertheless, no such "natural" example of a real number has been proven to be normal in *any* base, let alone absolutely normal. The conjectures that every algebraic irrational is absolutely normal and that π is absolutely normal are especially well known open problems [5, 7, 28].

This paper concerns an old problem, namely, the complexity of explicitly computing a real number that is provably absolutely normal, even if it is not natural in the above informal sense. Sierpinski and Lebesgue gave explicit constructions of absolutely normal numbers in 1917 [24, 17], but these were intricate limiting processes that offered no complexity analyses (coming two decades before the theory of computing) and little insight into the nature of the numbers constructed. In a 1936 note that was not published in his lifetime, Turing [26] gave a constructive proof that almost all real numbers are absolutely normal and then *derived* constructions of absolutely normal numbers from this proof. Moreover, although Turing does not mention Turing machines or computability, the note is typed, with equations handwritten by him, on the back of a draft of his paper on computable real numbers [25], so it is reasonable to interpret "constructively" in a computability-theoretic sense. And in fact his proof, with 2007 corrections by Becher, Figueira, and Picchi [1], explicitly computes an absolutely normal number. In 2013, Becher, Heiber, and Slaman [2] published an algorithm that computes an absolutely normal number in polynomial time. Specifically, this algorithm computes the binary expansion of an absolutely normal number x, with the nth bit of x appearing after $O(n^2 f(n))$ steps for any computable unbounded nondecreasing function f. (Unpublished polynomial-time algorithms for computing absolutely normal numbers were also announced independently by Mayordomo [21] and Figueira and Nies [12] at about the same time.)

In this paper we present a new algorithm that provably computes an absolutely normal in nearly linear time. Our algorithm computes the binary expansion of an absolutely normal number x, with the nth bit of x appearing after $O(n\text{polylog}(n))$ steps. The term "nearly linear time" was introduced by Gurevich and Shelah [13]. In that paper they showed that, while linear time computability is very model-dependent, nearly linear time is very robust. For example, they showed that random access machines, Kolmogorov-Uspensky machines, Schoenhage machines, and random-access Turing machines share exactly the same notion of nearly linear time.

Our algorithm uses the Lempel-Ziv parsing algorithm to achieve its nearly linear time bound. For each base $b \geq 2$, we use a martingale (betting strategy) that employs the Lempel-Ziv parsing algorithm and is implicit in the work of Feder [11]. This base-b Lempel-Ziv martingale succeeds exponentially when betting on the successive digits of the base-b expansion of any real number that is not normal in base b. Our algorithm simultaneously computes and diagonalizes against (limits the winnings of) a martingale that incorporates efficient proxies

of all these martingales, thereby efficiently computing a real number that is normal in every base.

The rest of this paper is organized as follows. Section 2 presents the base-b Lempel-Ziv martingales and their main properties. Section 3 shows how to transform a base-b Lempel-Ziv martingale into a base-b martingale with an efficiently computable nondecreasing savings account that is unbounded whenever the base-b Lempel-Ziv martingale succeeds exponentially. Section 4 develops an efficient method for converting a base-b martingale with an efficiently computable savings account to a base-2 martingale that succeeds whenever the base b savings account is unbounded. Section 5 presents an algorithm that exploits the uniformity of these constructions to efficiently and simultaneously compute (a) a single base-2 martingale d that succeeds on the binary expansion of every real number x for which some base-b martingale succeeds on the base-b expansion of x and (b) a particular real number x on which binary expansion d does not succeed. This x is, perforce, absolutely normal. Section 6 presents an open problem related to our work.

2 Lempel-Ziv Martingales

For each base $b \geq 1$ we let $\Sigma_b = \{0, 1, \ldots, b-1\}$ be the alphabet of *base-b digits*. We write Σ_b^* for the set of all (finite) *strings* over Σ_b and Σ_b^∞ for the set of all (infinite) *sequences* over Σ_b. We write $|x|$ for the length of a string or sequence x, and we write λ for the *empty string*, the string of length 0. For $x \in \Sigma_b^* \cup \Sigma_b^\infty$ and $0 \leq i \leq j < |x|$, we write $x[i..j]$ for the string consisting of the ith through jth digits in x. For $x \in \Sigma_b^* \cup \Sigma_b^\infty$ and $0 \leq n < |x|$, we write $x \upharpoonright n = x[0..n-1]$. For $w \in \Sigma_b^*$ and $x \in \Sigma_b^* \cup \Sigma_b^\infty$, we say that w is a *prefix* of x, and we write $w \sqsubseteq x$, if $x \upharpoonright |w| = w$.

A (*base-b*) *martingale* is a function $d : \Sigma_b^* \to [0, \infty)$ satisfying

$$d(w) = \frac{1}{b} \sum_{a \in \Sigma_b} d(wa) \tag{1}$$

for all $w \in \Sigma_b^*$. (This is the original martingale notion introduced by Ville [27] and implicit in earlier papers of Lévy [19, 20]. Its relationship to Doob's subsequent modifications [10], which are the "martingales" of probability theory, is explained in [15] along with the reason why Ville's original notion is still essential for algorithmic information theory.) Intuitively, a base-b martingale d is a *strategy for betting* on the subsequent digits in a sequence $S \in \Sigma_b^\infty$, with the strategy encoded in such a way that $d(S \upharpoonright n)$ is the amount of money that a gambler using the strategy d has after the first n bets. The condition (1) says that the payoffs for these bets are *fair* in the sense that the conditional expectation of $d(wa)$, given that w has occurred (and assuming that the digits $a \in \Sigma_b$ are equally likely), is $d(w)$.

A function $g : \Sigma_b^* \to [0, \infty)$ (which may or may not be a martingale) *succeeds*

on a sequence $S \in \Sigma_b^\infty$ if

$$\limsup_{n\to\infty} g(S \upharpoonright n) = \infty, \tag{2}$$

i.e., if its winnings on S are unbounded. The *success set* of a function $g : \Sigma_b^* \to [0,\infty)$ is

$$S^\infty[g] = \{S \in \Sigma_b^\infty \mid g \text{ succeeds on } S\}.$$

A function $g : \Sigma_b^* \to [0,\infty)$ *succeeds exponentially* on a sequence $S \in \Sigma_b^\infty$ if

$$\limsup_{n\to\infty} \frac{\log g(S \upharpoonright n)}{n} > 0, \tag{3}$$

i.e., if its winnings on S grow at some exponential rate, perhaps with recurrent setbacks. The *exponential success set* of a function $g : \Sigma_b^* \to [0,\infty)$ is

$$S^{\exp}[g] = \{S \in \Sigma_b^\infty \mid g \text{ succeeds exponentially on } S\}.$$

The $f(n)$ *success set* of a function $g : \Sigma_b^* \to [0,\infty)$ is

$$S^{f(n)}[g] = \left\{ S \in \Sigma_b^\infty \;\middle|\; \limsup_{n\to\infty} \frac{\log g(S \upharpoonright n)}{\log f(n)} \geq 1 \right\}.$$

Note that $S^{\exp}[g] = \cup_{\epsilon>0} S^{2^{\epsilon n}}[g]$.

A function $g : \Sigma_b^* \to [0,\infty)$ *succeeds strongly* on a sequence $S \in \Sigma_b^\infty$ if (2) holds with the limit superior replaced by a limit inferior i.e., if the winnings converge to ∞. A function $g : \Sigma_b^* \to [0,\infty)$ *succeeds strongly exponentially* on a sequence $S \in \Sigma_b^\infty$ if (3) holds with the limit superior replaced by a limit inferior i.e., if the winnings grow at exponential rate. The *strong success set* $S_{\text{str}}^{\exp}[g]$ of a function $g : \Sigma_b^* \to [0,\infty)$ are defined in the now-obvious manner. It is clear that the inclusions

$$S_{\text{str}}^{\exp}[g] \subseteq S^{\exp}[g] \subseteq S^\infty[g]$$

and

$$S_{\text{str}}^{\exp}[g] \subseteq S_{\text{str}}^\infty[g] \subseteq S^\infty[g]$$

hold for all $g : \Sigma_b^* \to [0,\infty)$.

For each base $b \geq 2$ the Lempel-Ziv martingale is a particular martingale $d_{\text{LZ}(b)}$ based on the Lempel-Ziv parsing algorithm [18], as we now explain.

Formally, $d_{\text{LZ}(b)}$ is computed by the algorithm in Figure 1, but some explanation is appropriate here. The algorithm is written with several instances of parallel assignment. For example, the second line initializes x, $L(x)$, and d to the values λ, 1, and 1, respectively. The items T, j, and $x(j)$ are not needed for the computation of $d_{\text{LZ}(b)}(w)$, but they are useful for understanding and analyzing the algorithm.

```
input w ∈ Σ_b^*;
x, L(x), d = λ, 1, 1;
T, j = {λ}, 0;
while true do
begin
    if w = λ then output d and halt;
    if L(x) = 1 then
    begin
        L(x) = b;
        for each 0 ≤ a < b do L(xa) = 1;
        T, x(j), j = T ∪ {x}Σ_b, x, j + 1;
        x = λ;
    end
    else
    begin
        a, w = head(w), tail(w);
        L(x), x, d = L(x) + b − 1, xa, (bL(xa)/L(x)) d
    end
end
```

Figure 1: Algorithm for computing $d_{\mathrm{LZ}(b)}(w)$.

The growing set T of strings in Σ_b^* always contains all the prefixes of all its elements, so it is a tree. We envision this tree as being oriented with its root at the top and the immediate children $v0, v1, \ldots, v(b-1)$ of each interior vertex v of T displayed left-to-right below T. The *dictionary* of the algorithm is the current set of leaves of T.

The string x in the algorithm is always an element of (i.e., location in) the tree T, and $L(x)$ is always the number of leaves of T that are descendants of x. We regard x as a descendant of itself, so x is a leaf if and only if $L(x) = 1$.

It is clear that $d_{\mathrm{LZ}(b)}(\lambda) = 1$. In fact, the algorithm's successive values of d are the values $d_{\mathrm{LZ}(b)}(u)$ for successive prefixes u of the input string w. More precisely, if w_t and d_t are the values of w and d after t executions of the else-block, then $w = (w \restriction t)w_t$ and $d_t = d_{\mathrm{LZ}(b)}(w \restriction t)$.

For $w \in \Sigma_b^*$ we define the tree $T(w)$ as follows. If $w = \lambda$, then $T(w) = \{\lambda\}$. If $w = w'a$, where $w' \in \Sigma_b^*$ and $a \in \Sigma_b$, then $T(wa)$ is the value of T when the algorithm terminates on input w'. (Note that this is one step before it terminates on input $w'a$.) For $w \in \Sigma_b^*$ we define $D(w)$ to be the number of leaves in $T(w)$.

The computation is divided into "epochs". At the beginning of each epoch, the string x is λ, i.e., it is located at the root of T. The string x then takes successive digits from whatever is left of w (because $a, w = \mathrm{head}(w), \mathrm{tail}(w)$ removes the first digit of w and stores it in a), following this path down the tree and updating d at each step, until w is empty (the end of the last epoch)

or x is a leaf if T. In the latter case, the jth epoch is over, the b children $x0, x1, \ldots, x(b-1)$ are added to T as new leaves, x is the jth *phrase* $x(j)$ of w, and x is reset to the root λ of T.

When the algorithm terminates, it is clear that exactly one of the following things must hold.

(a) $w = \lambda$.

(b) $w = x(1)\ldots x(j)$.

(c) $w = x(1)\ldots x(j)u$ for some nonempty interior vertex u of $T(w)$.

In case (a) or (b) we call w a *full parse*. In case (b) or (c) we call $x(1), \ldots, x(j)$ the *full phrases* of w. In case (c) we call u the *partial phrase* of w.

Define the set $A_b = \{1 + (b-1)r \,|\, r \in \mathbb{N}\}$ and the generalized factorial function $fact_b : A_b \to A_b$ by

$$fact_b(1 + (b-1)r) = \prod_{k=1}^{r} (1 + (b-1)k)$$

for all $r \in \mathbb{N}$. A routine argument, taking logarithms and bounding sums by integrals, gives the bounds

$$1 \leq \frac{fact_b(n)}{e^{\frac{1}{b-1}} \left(\frac{n}{e}\right)^{\frac{n}{b-1}}} \leq n \tag{4}$$

for all $n \in A_b$.

Lemma 2.1 *Let $w \in \Sigma_b^*$.*

1. *If w is a full parse, then*

$$d_{\mathrm{LZ}(b)}(w) = \frac{b^{|w|}}{fact_b(D(w))}.$$

2. *If w is not a full parse and u is its partial phrase, then*

$$d_{\mathrm{LZ}(b)}(w) = \frac{b^{|w|}}{fact_b(D(w))} L(u),$$

where $L(u)$ is the number of leaves below u in $T(w)$.

Lemma 2.2 *For $S \in \Sigma_b^\infty$ the following three conditions are equivalent.*

(a) $S \in S^{b^{(1-\alpha)n}}[d_{LZ(b)}]$.

(b) *There exist infinitely many full parses $w \sqsubseteq S$ for which*

$$D(w) \log_b |w| < \alpha(b-1)|w|.$$

(c) *There exist infinitely many full parses $w \sqsubseteq S$ for which*

$$D(w) \log_b D(w) < \alpha(b-1)|w|.$$

Corollary 2.3 *For $S \in \Sigma_b^\infty$ the following three conditions are equivalent.*

(a) $S \in S^{\exp}[d_{\mathrm{LZ}(b)}]$.

(b) *There exist $\alpha < 1$ and infinitely many full parses $w \sqsubseteq S$ for which*

$$D(w) \log_b |w| < \alpha(b-1)|w|.$$

(c) *There exist $\alpha < 1$ and infinitely many full parses $w \sqsubseteq S$ for which*

$$D(w) \log_b D(w) < \alpha(b-1)|w|.$$

We conclude this section by explaining the connection between the Lempel-Ziv martingales and normality. First, Schnorr and Stimm [23] defined (implicitly) a notion of *finite-state base-b martingale* and proved that every sequence $S \in \Sigma_b^\infty$ obeys the following dichotomy.

1. If S is normal, then no finite-state base-b martingale succeeds on S. (In fact, every finite-state base-b martingale decays exponentially on S.)

2. If S is not normal, then some finite-state base-b martingale succeeds exponentially on S.

Some twenty years later, Feder [11] defined (implicitely) the Lempel-Ziv martingale $d_{\mathrm{LZ}(b)}$ and proved (implicitly) that $d_{\mathrm{LZ}(b)}$ is at least as successful *on every sequence* as every finite-state base-b martingale. That is, for finite-state base-b martingale d, the inclusions

$$S^\infty[d] \subseteq S^\infty[d_{\mathrm{LZ}(b)}],\ S^\infty_{\mathrm{str}}[d] \subseteq S^\infty[d_{\mathrm{LZ}(b)}],$$

$$S^{\exp}[d] \subseteq S^{\exp}[d_{\mathrm{LZ}(b)}],\ S^{\exp}_{\mathrm{str}}[d] \subseteq S^{\exp}_{\mathrm{str}}[d_{\mathrm{LZ}(b)}]$$

all hold. This, together with Schnorr and Stimm's dichotomy result, implies that $d_{\mathrm{LZ}(b)}$ succeeds exponentially on every non-normal sequence in Σ_b^∞. Hence a real number x is absolutely normal if none of the martingales $d_{\mathrm{LZ}(b)}$ succeed exponentially on the base-b expansion of $x - \lfloor x \rfloor$.

3 Savings Accounts

In this section we will construct a conservative version of the the Lempel-Ziv martingale $d_{\mathrm{LZ}(b)}$, consisting of a new martingale d' that can be smaller than $d_{\mathrm{LZ}(b)}$ but that will have a savings account in the following sense.

Definition. A nondecreasing function $g : \Sigma_b^* \to [0, \infty)$ is a *savings account* of a martingale $d : \Sigma_b^* \to [0, \infty)$ if for every $w \in \Sigma_b^*$, $d(w) \geq g(w)$.

Construction 3.1 *Let $d = d_{LZ(b)}$ be the base-b-Lempel-Ziv martingale. Let $\alpha < 1$. We define a new martingale $d' = e' + g'$ as follows.*

We first define e'. $e'(\lambda) = b$. For $w \in \Sigma_b^*$, let $w = x(1)\dots x(j)u$, for $z = x(1)\dots x(j)$ a full parse and u a partial phrase of w.

$$goal(w) = |w| - \lceil D(z)(\log_b(D(z)))/(b-1)\rceil - 1$$

$$taken(w) = \begin{cases} taken(z) & \text{if } goal(w) \le \max\{taken(z), (1-\alpha|w|-2)\} \\ goal(w) & \text{if } goal(w) > \max\{taken(z), (1-\alpha|w|-2)\} \end{cases}$$

$$e'(w) = \begin{cases} e'(z)\cdot\frac{d(w)}{d(z)} & \text{if } goal(w) \le \max\{taken(z), (1-\alpha|w|-2)\} \\ e'(z)\cdot\frac{d(w)}{d(z)} b^{-goal(w)+taken(z)} & \text{if } goal(w) > \max\{taken(z), (1-\alpha|w|-2)\} \end{cases}$$

Let g' be defined as follows, $g'(\lambda) = 1$ and for $y \in \Sigma_b^*, a \in \Sigma_b$, if $ya = w$, let $w = x(1)\dots x(j)u$, for $z = x(1)\dots x(j)$ a full parse and u a partial phrase of w.

$$g'(ya) = \begin{cases} g'(y) & \text{if } goal(w) \le \max\{taken(z), (1-\alpha|w|-2)\} \\ g'(y) + e'(y)\frac{b-1}{b} & \text{if } goal(w) > \max\{taken(z), (1-\alpha|w|-2)\} \end{cases}$$

Claim 3.2 *$d' = e' + g'$ is a martingale and g' is a savings account for d'.*

Notice that g' is a monotonic function. Notice that in the second case $e'(ya) = e'(y)d(ya)/d(y)1/b$, so $\sum_{a\in\Sigma_b} e'(ya) + g'(ya) = e'(y)(1 + b\frac{b-1}{b}) + bg'(y) = b(e'(y) + g'(y))$.

Therefore g' is a savings account of d'.

Theorem 3.3 *Let d' be the martingale and g' its savings account defined in Construction 3.1. Then*

$$S^{b^{(1-\alpha)n}}[d_{LZ(b)}] \subseteq S^\infty[g'],$$

and d' and g' are both computable in online nearly linear time bounds that do not depend on b or α.

Proof of Theorem 3.3.

Claim 3.4 *If $w \in \Sigma_b^*$ is a full parse $w = x(1)\dots x(j)x(j+1)$ and $goal(w) \ge \max\{taken(x(1)\dots x(j)), (1-\alpha|w|-2)\}$ then $e'(w) = d(w)b^{-goal(w)}$.*

Claim 3.5 $S^{b^{(1-\alpha)n}}[d_{LZ(b)}] \subseteq S^\infty[g']$.

If $y \in S^{b^{(1-\alpha)n}}[d_{LZ(b)}]$ then by Lemma 2.2 for infinitely many n, $y \upharpoonright n$ is a full parse and $D(y\upharpoonright n)\log_b(D(y\upharpoonright n)) < \alpha(b-1)n$.

Notice that we use the second case in the definition of e' and g' and have that $g'(y\upharpoonright nb) = g'(y\upharpoonright n) + e'(y\upharpoonright n)\frac{b-1}{b}$.

By Claim 3.4 if $y \upharpoonright n = x(1)\dots x(j)x(j{+}1)$ and $goal(y \upharpoonright n) \geq \max\{taken(x(1)\dots x(j)), (1-\alpha n-2)\}$ then $e'(y \upharpoonright n) = d(y \upharpoonright n)b^{-goal(y\upharpoonright n)}$. Choose those (infinitely many) n for which, if $y \upharpoonright n = x(1)\dots x(j)x(j{+}1)$, $goal(y \upharpoonright n) \geq \max\{taken(x(1)\dots x(j)), (1-\alpha n-2)\}$, and $D(y \upharpoonright n)\log_b(D(y \upharpoonright n)) < \alpha(b-1)n$.

For those n, $e'(y \upharpoonright n) \geq b$, therefore $g'(y \upharpoonright nb) \geq g'(y \upharpoonright n) + b - 1$.

Since g' is monotonic, $y \in S^\infty[g']$.

Claim 3.6 *d' can be computed in online nearly linear time (in time $n\log^c n$ for c not depending on b or α).*

$w = x(1)\dots x(j)u$, let $z = x(1)\dots x(j)$. If $goal(w) > taken(z)$, let t be such that $t \sqsubseteq u$ and $goal(zt) = taken(z)$.

Notice that

$$g'(w) = \begin{cases} g'(z) & \text{if } goal(w) \leq \max\{taken(z), (1-\alpha|w|-2)\} \\ g'(z) + \sum_{v\sqsubseteq u, |zv|>|zt|} e'(zt)\frac{b-1}{b}L(v) & \text{if } goal(w) > \max\{taken(z), (1-\alpha|w|-2)\} \end{cases}$$

Where $L(v)$ is the number of leaves below v in $T(w)$. Given precomputed values for $f(u) = \sum_{v\sqsubseteq u} L(v)$, the value of $g'(w)$ can be easily computed in nearly linear time. The precomputed value of $f(u)$ will be stored in the u node of the LZ tree for z.

Notice that the whole computation can be done online.

□

4 Base change

We will use infinite sequences over Σ_b to represent real numbers in [0,1). For this, we associate each string $w \in \Sigma_b^*$ with the half-open interval $[w]_b$ defined by $[w]_b = [x, x + b^{-|w|})$, for $x = \sum_{i=1}^{|w|} w[i-1]b^{-i}$. Each real number $\alpha \in [0,1)$ is then represented by the unique sequence $\mathrm{seq}_b(\alpha) \in \Sigma_b^\infty$ satisfying

$$w \sqsubseteq \mathrm{seq}_b(\alpha) \iff \alpha \in [w]_b$$

for all $w \in \Sigma_b^*$. We have

$$\alpha = \sum_{i=1}^{\infty} \mathrm{seq}_b(\alpha)[i-1]b^{-i}$$

and the mapping $\alpha \mapsto \mathrm{seq}_b(\alpha)$ is a bijection from $[0,1)$ to Σ_b^∞ (notice that $[w]_b$ being half-open prevents double representations). If $x \in \Sigma_b^\infty$ then $\mathrm{real}_b(x) = \alpha$ such that $x = \mathrm{seq}_b(\alpha)$. Therefore we always have that $\mathrm{real}_b(\mathrm{seq}_b(\alpha)) = \alpha$ and $\mathrm{seq}_b(\mathrm{real}_b(x)) = x$. A set of real numbers $A \subseteq [0,1)$ is represented by the set

$$\mathrm{seq}_b(A) = \{\mathrm{seq}_b(\alpha) \mid \alpha \in A\}$$

of sequences. If $X \subseteq \Sigma_b^\infty$ then

$$\mathrm{real}_b(X) = \{\mathrm{real}_b(x) \mid x \in X\}.$$

For a positive number x we will use the notation $\log_b(x) = \lceil \log_b(x) \rceil$.

Construction 4.1 *Let $d : \Sigma_b^* \to [0, \infty)$ be a martingale with a savings account g.*

We define $\gamma : \Sigma_b^* \to [0, 1]$ a probability measure $\gamma(w) := b^{-|w|} d(w)$.

Using Carathéodory extension to Borel sets, γ can be extended to any interval $[a, c]$, we denote with $\hat{\gamma}$ this extension. (In fact if we consider all $U \subseteq \Sigma_b^*$ such that all $u, v \in U, u \neq v$ are incomparable and $[u]_b \subseteq [a, c]$ for all $u \in U$, then $\hat{\gamma}([a, c]) = \sup_U \sum_{u \in U} \gamma(u)$).

We define $\mu : \{0, 1\}^* \to [0, 1]$ as $\mu(y) = \hat{\gamma}([y]_2)$.

Finally we define $d^{(2)} : \{0, 1\}^* \to [0, \infty)$ as $d^{(2)}(y) = 2^{|y|}\mu(y)$.

Theorem 4.2 *Assume that d is a base-b martingale and g is a savings account of d, ant let $d^{(2)}$ be defined from d and g as in Construction 4.1. Then*

$$real_b(S^\infty[g]) \subseteq real_2(S^\infty[d^{(2)}]).$$

Moreover, if d and g are computable in online nearly linear time bounds not depending on b, then so is $d^{(2)}$.

Property 4.3 *Let $\alpha \in [0, 1)$. If $\mathrm{seq}_b(\alpha) \in S^\infty[g]$ then $\mathrm{seq}_2(\alpha) \in S^\infty[d^{(2)}]$.*

Proof of Property 4.3. Let $x = \mathrm{seq}_b(\alpha) \in S^\infty[g]$. Let $y = \mathrm{seq}_2(\alpha)$.

We use here that d has a savings account g, so if $g(x \upharpoonright n) > m$ then for all w with $x \upharpoonright n \sqsubseteq w$, $g(w) > m$.

Let $m \in \mathbb{N}$ then there is an n such that $g(x \upharpoonright n) > m$. Let q be such that $[y \upharpoonright q]_2 \subseteq [x \upharpoonright n]_b$. Let us see that $d^{(2)}(y \upharpoonright q) > m$.

Let $r \in \mathbb{N}$. Let $A_r^q = \{w \in \Sigma_b^* \mid |w| = r \text{ and } [w]_b \subseteq [y \upharpoonright q]_2\}$.

$d^{(2)}(y \upharpoonright q) = 2^q \hat{\gamma}([y \upharpoonright q]_2) = 2^q \lim_r \sum_{w \in A_r^q} d(w) b^{-|w|} \geq 2^q m \lim_r \sum_{w \in A_r^q} b^{-|w|} = 2^q m 2^{-q} = m$.

The last chain of equations is based on $[y \upharpoonright q]_2 \subseteq [x \upharpoonright n]_b$ and for every $w \in A_r^q$, $[w]_b \subseteq [y \upharpoonright q]_2$, therefore $x \upharpoonright n \sqsubseteq w$ for any $w \in A_r^q$.

□

For each $m \in \mathbb{N}$ we define $\mu_m : \{0, 1\}^* \to [0, 1]$ as follows

$$\mu_m(y) = \sum_{|w|=m, [w]_b \cap [y]_2 \neq 0} \gamma(w)$$

Property 4.4 *For every $y \in \{0, 1\}^*$, for every $m \in \mathbb{N}$,*

$$|\mu(y) - \mu_m(y)| \leq 2b^{-m}(b^3 m + d(\lambda)).$$

For each $m \in \mathbb{N}$ we define $d_m^{(2)} : \{0, 1\}^* \to [0, \infty)$ as follows $d_m^{(2)}(y) = 2^{|y|}\mu_m(y)$.

Property 4.5 *For some $c > 0$, for every $y \in \{0,1\}^*$, for every $m \in \mathbb{N}$,*

$$|d^{(2)}(y) - d^{(2)}_m(y)| \leq 2^{|y|} 2^{-m \log b + \log m + c}.$$

Corollary 4.6 *For every $y \in \{0,1\}^*$, for every $m \in \mathbb{N}$,*

$$|d^{(2)}(y) - d^{(2)}_{|y|/\log b + 4 \log |y|} d(y)| \leq 1/|y|^3.$$

Notice that the approximation of $d^{(2)}$ is slower than usual, but it is enough for the diagonalization performed in the next section.

Property 4.7 *For $m \in \mathbb{N}$, $d^{(2)}_m(y)$ can be computed by considering a maximum of 4 neighbor strings $w \in \Sigma^r_b$ for $r = |y|/\log b$ to m, computing $d(w)$ for each of them and doing an addition and a multiplication for each.*

Corollary 4.8 *For $m \in \mathbb{N}$, $d^{(2)}_{|y|/\log b + 4 \log |y|}(y)$ can be computed by considering a maximum of 4 neighbor strings $w \in \Sigma^r_b$ for $r = |y|/\log b$ to $|y|/\log b + 4 \log |y|$, computing $d(w)$ for each of them and doing an addition and a multiplication for each.*

This concludes the proof of Theorem 4.2.

5 Absolutely Normal Numbers

In this section we will give an algorithm that will diagonalize on the Lempel-Ziv martingales for all bases and that will work in nearly linear time.

We will use the following theorem

Theorem 5.1 *Let $(d_k)_{k \in \mathbb{N}}$ be a sequence of base-2 martingales such that for each of them there exists a function $\widehat{d_k} : \{0,1\}^* \to [0, \infty)$ with the following two properties*

1. *for every $y \in \{0,1\}^*$*
$$|d_k(y) - \widehat{d_k}(y)| \leq \frac{1}{|y|^3}$$
2. *$\widehat{d_k}$ is computable in online nearly linear time bounds that do not depend on k.*

then we can compute in nearly linear time a binary sequence x such that for every k, $x \notin S^\infty[d_k]$.

Proof. W.l.o.g. we assume that $d_k(\lambda) = 1$ for all k.

Let $d : \{0,1\}^* \to [0, \infty)$ be defined as

$$d(w) = \sum_{k=1}^{\log^2 |w|} 2^{-k - 2^{\sqrt{k}}} \widehat{d_k}(w).$$

Our algorithm will "diagonalize" against d constructing a binary sequence x as follows, if $x \upharpoonright n$ has been defined then choose the next bit of x as $c \in \{0,1\}$ that minimizes $d((x \upharpoonright n)c)$.

Claim 5.2 *If $x \notin S^\infty[d]$ then for every k, $x \notin S^\infty[d_k]$.*

Let $n \in \mathbb{N}, k \in \mathbb{N}$, $n \geq 2^{\sqrt{k}}$.

$$d_k(x \upharpoonright n) \leq \widehat{d_k}(x \upharpoonright n) + 1/n^3 \leq$$

$$\leq 2^{k+2^{\sqrt{k}}} d(x \upharpoonright n) + 1/n^3$$

Claim 5.3 *$x \notin S^\infty[d]$.*

Let n be such that $\log^2 n = \log^2(n-1) + 1$.

$$d(x \upharpoonright n) \leq 2^{-\log^2 n} + 1/n^3 + \sum_{k=1}^{\log^2(n-1)} 2^{-k-2^{\sqrt{k}}} \widehat{d_k}(x \upharpoonright (n-1)) + 2/n^3 = 2^{-\log^2 n} + 3/n^3 + d(x \upharpoonright (n-1))$$

If $\log^2 n = \log^2(n-1)$ then

$$d(x \upharpoonright n) \leq \sum_{k=1}^{\log^2(n-1)} 2^{-k-2^{\sqrt{k}}} \widehat{d_k}(x \upharpoonright (n-1)) + 2/n^3 = 2/n^3 + d(x \upharpoonright (n-1))$$

□

6 Open Problem

Many questions arise naturally from this work, but the following problem appears to be especially likely to demand new and useful methods.

As we have seen, normal numbers are closely connected to the theory of finite automata. Schnorr and Stimm [23] proved that normality is exactly the finite-state case of randomness. That is, a real number α is normal in a base $b \geq 2$ if and only if no finite-state automaton can make unbounded money betting on the successive digits of the base-b expansion of α with fair payoffs. The theory of finite-state dimension [9], which constrains Hausdorff dimension [14] to finite-state automata, assigns each real number α a finite-state dimension $\dim_{\mathrm{FS}}^{(b)}(\alpha) \in [0,1]$ in each base b. A real number α then turns out to be normal in base b if and only if $\dim_{\mathrm{FS}}^{(b)}(\alpha) = 1$ [6]. Do there exist *absolutely dimensioned numbers*, i.e., real numbers α for which $\dim_{\mathrm{FS}}(\alpha) = \dim_{\mathrm{FS}}^{(b)}(\alpha)$ does not depend on b, and $0 < \dim_{\mathrm{FS}}(\alpha) < 1$?

Acknowledgments

The first author's research was supported in part by National Science Foundation Grants 0652569, 1143830, 1247051, and 1545028. Part of this author's work was done during a sabbatical at Caltech and the Isaac Newton Institute for Mathematical Sciences at the University of Cambridge, part was done during three weeks at Heidelberg University, with support from the Mathematics Center Heidelberg and the Heidelberg University Institute of Computer Science, and part was done during the workshop "Normal Numbers: Arithmetical, Computational and Probabilistic Aspects" at the Erwin Schrödinger International Institute for Mathematics and Physics at the University of Vienna.

The second author's research was supported in part by Spanish Government MEC Grants TIN2011-27479-C04-01 and TIN2016-80347-R. Part of this author's work was done during a research stay at the Isaac Newton Institute for Mathematical Sciences at the University of Cambridge, part was done during three weeks at Heidelberg University, with support from the Mathematics Center Heidelberg and the Heidelberg University Institute of Computer Science, and part was done during the workshop "Normal Numbers: Arithmetical, Computational and Probabilistic Aspects" at the Erwin Schrödinger International Institute for Mathematics and Physics at the University of Vienna.

References

[1] V. Becher, S. Figueira, and R. Picchi. Turing's unpublished algorithm for normal numbers. *Theoretical Computer Science*, 377:126–138, 2007.

[2] V. Becher, P. A. Heiber, and T. A. Slaman. A polynomial-time algorithm for computing absolutely normal numbers. Manuscript, 2013.

[3] P. Billingsley. *Probability and Measure,* third edition. John Wiley and Sons, New York, N.Y., 1995.

[4] E. Borel. Sur les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo*, 27(1):247–271, 1909.

[5] J. Borwein and D. Bailey. *Mathematics by Experiment: Plausible Reasoning in the 21st Century,* second edition. A. K. Peters, 2008.

[6] C. Bourke, J. M. Hitchcock, and N. V. Vinodchandran. Entropy rates and finite-state dimension. *Theoretical Computer Science*, 349(3):392–406, 2005.

[7] Y. Bugeaud. *Distribution modulo one and Diophantine approximation*, volume 193. Cambridge University Press, 2012.

[8] J. W. S. Cassels. On a problem of Steinhaus about normal numbers. *Colloquium Mathematicum*, 7:95–101, 1959.

[9] J. J. Dai, J. I. Lathrop, J. H. Lutz, and E. Mayordomo. Finite-state dimension. *Theoretical Computer Science*, 310:1–33, 2004.

[10] J. Doob. Regularity properties of certain families of chance variables. *Transaction of the American Mathematical Society*, 47:455–486, 1940.

[11] M. Feder. Gambling using a finite state machine. *IEEE Transactions on Information Theory*, 37:1459–1461, 1991.

[12] S. Figueira and A. Nies. Feasible analysis and randomness. Manuscript, 2013.

[13] Y. Gurevich and S. Shelah. Nearly linear time. *Springer Lecture Notes in Computer Science*, 363:108–118, 1989.

[14] F. Hausdorff. Dimension und äußeres Maß. *Math. Ann.*, 79:157–179, 1919.

[15] J. Hitchcock and J. Lutz. Why computational complexity requires stricter martingales. *Theory of Computing Systems*, 39:277–296, 2006.

[16] A. N. Kolmogorov. *Foundations of the Theory of Probability*. Chelsea, 1950.

[17] H. Lebesgue. Sur certaines demonstrations d'existence. *Bull. Soc. Math. de France*, 45:132–144, 1917.

[18] A. Lempel and J. Ziv. Compression of individual sequences via variable rate coding. *IEEE Transaction on Information Theory*, 24:530–536, 1978.

[19] P. Lévy. Propriétés asymptotiques des sommes de variables indépendantes ou enchainées. *Journal des mathématiques pures et appliquées. Series 9.*, 14(4):347–402, 1935.

[20] P. Lévy. *Théorie de l'Addition des Variables Aleatoires*. Gauthier-Villars, 1937 (second edition 1954).

[21] E. Mayordomo. Construction of an absolutely normal real number in polynomial time. Manuscript, 2013.

[22] W. M. Schmidt. On normal numbers. *Pacific J. Math*, 10(2):661–672, 1960.

[23] C.-P. Schnorr and H. Stimm. Endliche automaten und zufallsfolgen. *Acta Informatica*, 1(4):345–359, 1972.

[24] W. Sierpinski. Démonstration élémentaire du théorème de M. borel sur les nombres absolument normaux et détermination effective d'une tel nombre. *Bull. Soc. Math. France*, 45:125–132, 1917.

[25] A. M. Turing. On computable numbers with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265, 1936-1937.

[26] A. M. Turing. A note on normal numbers. In S. B. Cooper and J. van Leeuwen, editors, *Alan Turing: His Work and Impact*. Elsevier Science, 2013.

[27] J. Ville. *Étude Critique de la Notion de Collectif*. Gauthier–Villars, Paris, 1939.

[28] S. Wagon. Is π normal? *Math. Intelligencer*, 7:65–67, 1985.